

AD-A212 543

Investigation of Adaptive Protocols
for Spread-Spectrum Packet Radio Networks

FINAL REPORT

M. B. Pursley and H. B. Russell

SIGCOM, Inc.
1776 E. Washington
Urbana, IL 61801

July 31, 1989

DTIC
ELECTE
SEP. 13 1989
S B D

Contract Number: DAAL03-88-C-0022

The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U. S. Army Research Office or the U. S. Government.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S) ARL 25947.2-EL-SBI	
6a. NAME OF PERFORMING ORGANIZATION SIGCOM, Inc.	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION U.S. Army Research Office	
6c. ADDRESS (City, State, and ZIP Code) 1776 East Washington Urbana, IL 61801		7b. ADDRESS (City, State, and ZIP Code) P.O. Box 12211 Research Triangle Park, NC 27709	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U.S. Army Research Office	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DAAL03-88-C-0022	
8c. ADDRESS (City, State, and ZIP Code) P.O. Box 12211 Research Triangle Park, NC 27709		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) Investigation of Adaptive Protocols for Spread-Spectrum Packet Radio Networks			
12. PERSONAL AUTHOR(S) M. B. Pursley and H. B. Russell			
13a. TYPE OF REPORT Final	13b. TIME COVERED FROM 88-9-1 TO 89-7-31	14. DATE OF REPORT (Year, Month, Day) 89-7-31	15. PAGE COUNT 71
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
		Frequency-hop packet radio networks, adaptive protocols	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This report describes research in adaptive protocols for spread-spectrum packet radio networks. Adaptive forwarding and routing protocols were developed and their performance was determined by simulation of a frequency-hop packet radio network with mobile partial-band jamming. New forwarding protocols were designed to permit short-term adaptation to changes in the interference environment in the network. A new routing technique, called <i>least-resistance routing</i> was developed, and various versions of this routing method were examined. Comparisons were made between different versions of least-resistance routing and between least-resistance routing and previously developed adaptive routing techniques. It was found that least-resistance routing is an effective way for dealing with mobile jamming in a frequency-hop packet radio network. Significant increases in throughput and end-to-end probability of success were obtained by use of least-resistance routing.			
20. SECURITY CLASSIFICATION <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL		22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL

FINAL REPORT

Investigation of Adaptive Protocols for Spread-Spectrum Packet Radio Networks

prepared by
M. B. Pursley and H. B. Russell

SIGCOM, Inc.
1776 E. Washington
Urbana, IL 61801

SUMMARY

This report describes research in adaptive protocols for spread-spectrum packet radio networks. Adaptive forwarding and routing protocols were developed and their performance was determined by simulation of a frequency-hop packet radio network with mobile partial-band jamming. New forwarding protocols were designed to permit short-term adaptation to changes in the interference environment in the network. A new routing technique, called *least-resistance routing* was developed, and various versions of this routing method were examined. Comparisons were made between different versions of least-resistance routing and between least-resistance routing and previously developed adaptive routing techniques. It was found that least-resistance routing is an effective way for dealing with mobile jamming in a frequency-hop packet radio network. Significant increases in throughput and end-to-end probability of success were obtained by use of least-resistance routing.

The research described in this report was supported by the U.S. Army Research Office under Contract Number DAAL03-88-C-0022. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U. S. Army Research Office or the U. S. Government.

1. Introduction

New methods for adaptive forwarding and routing in a spread-spectrum packet radio network with localized jamming have been devised and tested as part of the research completed during Phase I. The research has concentrated on frequency-hop (FH) packet radio networks, and the sources of interference of primary interest are other FH radios that are part of the network, FH radios that are not part of the network (e.g., voice radios), and mobile partial-band jamming.

The purpose of the *forwarding protocol* is to permit a radio to make changes to its normal routing procedures when congestion or partial-band jamming is introduced into a portion of the network. As a result, for the work in forwarding protocols, the primary interest is in short-term changes, and the simulation that we have used to study forwarding protocols concentrates on the time period between routing updates. The changes in routing made by the forwarding protocol affect only routes near the interference or congestion, and they require local information only in order to determine what changes to make. In particular, the forwarding protocols that we have designed and investigated permit the radio to adjust its local short-term routing based only on acknowledgements from its neighbors.

The purpose of the *routing protocol* is to make long-term changes to the routes between source-destination pairs in the network. Unlike the forwarding protocol, the routing protocol bases its decisions on information from radios throughout the network. The information of primary interest concerns the interference environment at the individual radios, and this is passed around the network on special organization packets and on the data packets. The routing protocol developed and simulated in this project is an adaptive, decentralized, distributed routing protocol with a metric that is specifically designed for the interference that will be experienced by a FH radio.

on For	
RA&I	<input checked="" type="checkbox"/>
3	<input type="checkbox"/>
need	<input type="checkbox"/>
station	

By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



The radios in the network use time-slotted, receiver-directed, FH spread-spectrum signaling, which provides both multiple-access capability and anti-jam communications. We consider different types of FH radios: those that are interacting closely in a "local" network that we refer to as a *subnetwork*, those that are FH packet radios but are not part of the subnetwork (although they may be able to reach the subnetwork via gateways), and other FH radios. All of these radios contribute interference in the subnetwork, but only the radios within the subnetwork interact in the routing protocol.

A computer program has been developed to simulate the subnetwork. The simulation models the processing and transmission of all packets for each radio in the subnetwork. A combination of simulation and analysis is used to account for the influence of other FH radios in the network as well as FH radios that are not part of the network. Certain of the packets are *tagged*, and their flow through the subnetwork is monitored. The results on throughput, delay, and success probability that we obtained are based exclusively on the tagged packets. The other packets that enter the subnetwork represent sources of interference only. The origination of these packets is simulated, but their effects on the tagged packets are handled analytically, as are the effects of transmissions by FH radios outside the network. For each link in the subnetwork, we obtain from the simulation the total number of interfering transmissions in a particular packet interval. From this the success probabilities on the individual links for that interval are calculated, and these probabilities are inputs to the simulation program.

One key feature of this simulation is that it uses a realistic model of the interference environment for an tactical FH radio network [10]. The interference due to other transmissions and partial-band jamming is accounted for in the computation of the success probabilities for the links, the two components of which

are the probability of decoding the packet correctly and the probability of acquiring synchronism.

We consider static network topologies only, but dynamic network jamming is included. The purpose of the simulation is to examine the network's ability to route around jammed nodes with a minimum of information exchange among the radios. In our simulation, the only information available to a radio concerning the fate of the transmission of a packet to one of its neighbors is that which can be obtained from acknowledgements: either a packet is acknowledged or not. Primary and secondary routes from a given radio to a given destination are determined by the routing protocol, but each radio along the route will decide, by means of its forwarding protocol, whether to use the primary route or the secondary route. The decision made by a given radio is based on the outcomes of previous transmission attempts by that radio.

2. Overview of Research Activity

SIGCOM has continued to develop and improve its FH packet radio network simulation during the Phase I research program that began on September 1, 1988. This simulation models multiple-access interference from FH packet radio networks and FH voice radios. It also models stationary and mobile partial-band jamming. We have employed the simulation to design and test the adaptive protocols that were developed in the Phase I program.

For purposes of comparison and as a first step in implementing other routing protocols, we implemented the routing protocol known as *tier routing* [5,6] in our simulation. The actual version of tier routing that we implemented was suitably modified for a FH packet radio network, and it includes some of the features we

found important from our investigations of forwarding protocols. In tier routing, as in our other adaptive routing techniques, each radio periodically broadcasts a packet radio organization packet (PROP) [1,6]. The PROP's are typically transmitted on a common spread-spectrum waveform so that all radios that are within range of a given PROP transmission can receive it. From a neighbor's PROP, a radio determines certain statistics regarding successes and failures of packet transmissions to and from that neighbor. Based on these statistics, the neighbor is flagged as good or bad, and this status is stored in the radio's neighbor table. Each radio maintains a routing table with an entry for each destination. For a given destination, the entry includes a primary and secondary outgoing link for that destination (for use in our forwarding protocols) and the number of hops on each route to the destination. Each outgoing link listed must be to a good neighbor. When a radio transmits a PROP, it includes a designation of the primary outgoing link and the number of hops to each destination. When a radio receives a PROP from a good neighbor, it updates its routing table in the event that, for a particular destination, the neighbor (i) provides a new route with fewer hops, (ii) no longer provides a route, or (iii) has a new route length for an existing route. In addition, when a radio receives a packet for a destination that it no longer has a routing entry for, it includes a flag in its acknowledgement so the neighbor will clear the incorrect routing entry. We have made simulation runs to provide the necessary data for investigating the performance of the network while using tier routing in conjunction with the various forwarding protocols that we have designed in the Phase I investigation.

We also used the tier routing formulation as a mechanism for the implementation of least-resistance routing (LRR), a new adaptive routing method developed in our Phase I investigation. The LRR protocol uses a quantitative assessment of the interference environment at each radio along a given route to determine the resistance that is assigned to that route. This necessitates the

estimation of the interference that will be encountered at each receiver in the network. For the model considered in our investigation, the interference has two components: transmissions from other radios that are within range of the receiver and partial-band jamming. Our first implementation assumed that each radio knows the number of transmissions that are within range and the number of frequency slots that are jammed. During the Phase I program, we have examined some aspects of the problem of determining the interference environment in the receiver of a frequency-hop radio. The second implementation assumes only that the radio can measure the number of frequency slots that contain some type of interference. In this approach, discussed in more detail in Sections 4 and 5, a receiver estimates the expected number of frequency slots that contain some type of interference. This is used to obtain an estimate of that receiver's success probability (i.e., the probability that a packet can be acquired and decoded at the radio in question).

We have defined the *resistance* for a link to be the negative of the logarithm of the success probability for that link. The *path resistance* is the sum of the resistance values for the links that make up the path, so, with our scheme, the path resistance is equal to the negative of the logarithm of the product of the success probabilities for the link's in the path. Of course, successes on the links in a path are not statistically independent, so no claim is made that our link resistance is a measure of the overall success probability for the path. In addition, several other factors contribute to the success or failure of a packet transmissions. However, our results indicate that the proposed method for calculating path resistance gives a good measure of the interference along the routes in the network. For a network using the LRR protocol, a radio will first calculate the resistance for each of its neighbors based on information contained in the PROP's. Routes are then selected

on the basis of the total resistance to a particular destination, rather than on the minimum number of links to the destination as in tier routing.

After the simulation program was enhanced to include the option of using either LRR or tier routing, we selected four sets of simulations to explore the behavior of LRR and make comparisons with tier routing. The first set of simulations permits a comparison of LRR and tier routing in the presence of a *stationary* jammer. Two jamming scenarios were considered, one with 55% of the band jammed and the other with 40% of the band jammed. For this and the second set of simulations, it was assumed that the FH radios have perfect knowledge of the interference environment, and, therefore, were able to determine exactly the link parameters required to compute the link resistances. The second set of simulations permits a comparison of the performance of LRR and tier routing in the presence of a *mobile* jammer. We considered a jammer that moves once every 1000 time intervals (a time interval is the time required to transmit one packet) in a network for which PROP's are transmitted every 240 time intervals. We also investigated jammers which move more slowly and radio networks which transmit PROP's more frequently. In the third set of simulations, the effects of errors in the estimates of interference levels on the performance of LRR were examined (i.e., a sensitivity analysis). The fourth set of simulations explored some improvements to the basic LRR algorithms and to the network protocols. Further details, including graphs of simulation results, are presented in Section 6.

During Phase I, we increased the number of radios in our simulation from eight to twelve. This provides a better topology for examining the performance of the routing algorithms. In particular, the larger subnetwork has the ability to develop short routes with higher resistance than some of the longer routes. This is especially important for comparing LRR with shortest-path routing methods such as tier routing.

We have been invited to present a paper at MILCOM '89 based on our work on adaptive forwarding and routing in frequency-hop packet radio networks [18]. A manuscript was approved by the Army Research Office for presentation at this conference.

3. Models and Protocols Employed in the FH Packet Radio Network Simulation

Our model for the FH packet radio network is illustrated in Figure 3.1. Two different *subnetworks* have been examined; the first subnetwork consists of the radios numbered 1-8 in Figure 3.1, and the second subnetwork consists of the radios numbered 1-12.




Figure 3.1 Model for the packet radio network topology

The computer program actually simulates the generation, processing, and transmission of all packets for each radio in the *subnetwork*. A combination of simulation and analysis is used to account for the influence of other FH packet radios. To measure the performance of the network and its protocols, certain packets are monitored as they travel through the network. The end-to-end throughput, delay, and probability of success for these packets are calculated. The other packets that enter the network represent sources of interference only. These packets must be stored and forwarded by the radios in the subnetwork, but these packets are not included in the performance measures.

The simulation uses a technique that employs the regenerative properties of networks of queues to permit the calculation of confidence intervals. This approach overcomes problems due to the statistical dependence between data points, the initial transients in the queue behavior, and the packets still in the network when the simulation is stopped.

For receiver-directed FH signaling, a radio must transmit on the hopping pattern of the intended recipient. This requires the use of *active* acknowledgements. In a time-slotted network, each data packet and acknowledgement packet require an entire packet interval, and this results in increased delay as compared with a broadcast network that employs passive acknowledgements. On the other hand, the use of receiver-directed transmission holds the promise of increased throughput because of the multiple-access capability it provides.

In our simulation, acknowledgements are given priority over all other packets. After sending a packet, the transmitting radio switches to the receive mode to look for an acknowledgement packet in the next time interval. After receiving a packet, the receiving radio switches immediately to the transmit mode and sends an acknowledgement packet in the next time slot. Prompt acknowledgement of packets is important in this system since single threading [1] is employed.

In the development of the LRR algorithm, we found it beneficial to include some of the techniques of tier routing to help avoid generating routes that contain loops. Specifically, for each destination, each radio not only stores the total resistance to that destination, but it also keeps track of the number of hops to that destination. Both the number of hops to each destination and the resistance values are included in the PROP. When a PROP is received, the LRR algorithm examines the change in the length of each route for evidence of a loop. Also, when a packet is forwarded, the number of hops to the destination is included in the packet's header so that the progress of the packet can be monitored to determine if the packet is traveling in a loop.

In addition to adjusting the routing tables whenever a PROP is received, the resistance values are updated each time a data packet is received. An extra field has been included in the header of each packet for the resistance at the transmitting radio. When the packet is received, the resistances for all the routes at the receiving radio that pass through the transmitting radio are updated to include the new resistance. Also, when a packet is acknowledged, the resistance between the acknowledging radio and the packet's destination is included in the acknowledgement, so that the resistance for that route can be updated at the radio that receives the acknowledgement. Thus, after a few packets have been forwarded along a given route, resistance data is reflected from the destination of the route back to the source.

The resistances are used by the routing algorithm to determine a primary and secondary route from each origin to each destination. Some of the forwarding protocols that we investigated adapt to the resistance values on the available routes. The basic forwarding algorithm uses the primary route for the first forwarding attempts and, if unsuccessful, it employs the secondary route for the remaining attempts. When two routes have acceptable resistance values, the adaptive

forwarding algorithm alternates packets on both routes. The amount of traffic on a particular route will depend on the relative resistance of that route. A forwarding protocol was implemented to alternate packets between the primary and secondary routes. The number of packets selected to be forwarded on the secondary route depends on the relative resistance between the primary and secondary routes. The *n*-packet forwarding protocol forwards every *n*th packet on the secondary instead of the primary route. The value of *n* depends on the difference in the path resistance for the primary and secondary routes.

All of our results to date show that LRR adapts more quickly than tier routing to a change in the position of a mobile jammer. This is as expected, because the LRR algorithm is not limited by having to wait for a PROP to update its routing tables, since each radio updates some of its resistance values after every reception. The resistance of a link to a radio is updated during every time interval in which the radio is not transmitting (the radio must be in the receive mode in order to estimate the interference level). Moreover, the resistance values for the a link provide more information about the chances for success of a transmission on that link than is provided in tier routing. Tier routing, on the other hand, does not react quickly to a mobile jammer, because the next update cannot take place until the next PROP is received. Also, in tier routing the statistics required to determine the status of a link are obtained from packet transmissions and receptions that take place on this link. As a result, the statistics that are determined are not reliable unless there have been a fairly large number of transmissions on the link since the last change in the vicinity of this link (i.e., since the introduction of partial-band jamming).

One of the issues that arises in our simulation and in a fielded packet radio network is the separation of multiple-access interference from other forms of partial-band interference. The former will be much more dynamic than the latter, and it has a different effect on radio performance. For present radios, such as

SINCGARS, we feel it is unlikely that it will be possible to distinguish between these two types of interference. For future radios, however, it may be possible to include detection subsystems that will permit a distinction between multiple-access interference and the more nearly stationary partial-band interference. The assumption made in our recent simulations is that a radio can determine an estimate of the number of frequency slots with interference, without being able to determine how many of these are due to multiple-access interference and how many are due to partial-band jamming or other forms of partial-band interference.

Examining the performance of the routing algorithm closely, we found that the resistance value at a radio increases quickly when a jammer moves into its neighborhood, but when a jammer moves away from a radio the neighboring radios do not discover this very quickly. Consequently, after a jammer moves new routes using the radio that is no longer jammed are not discovered very quickly. To allow these new routes to be discovered more quickly, an *event-driven* PROP protocol was introduced. When the resistance at a radio changes significantly, the radio generates a PROP transmission.

The simulation model of the interference environment was improved so that PROP transmissions are subject to interference (all other transmissions have always been subject to the interference model). We have given PROP transmissions priority over all other transmissions by reserving the first three dwell intervals of every time interval for the synchronization dwell intervals of the PROP transmissions. A radio will first attempt to acquire synchronism with a PROP transmission and if it fails it will try to acquire synchronism with some other transmission on its hopping pattern. The radio can also fail to receive a PROP if the packet cannot be decoded. Of course the radio can not attempt to receive another packet in this case.

4. Least-Resistance Routing: A New Adaptive Routing Protocol for FH Packet Radio Networks

Least-resistance routing (LRR) is an adaptive, decentralized routing algorithm that accounts for multiple-access interference, jamming, and other partial-band interference at each of the radios in the network. Each radio maintains a measure of its own reception quality for packets coming from each of its neighbors, and it can pass this information along to other radios by means of a number of different mechanisms including PROP's and data packets. The metric for least-resistance routing accurately reflects the channel conditions *as seen by a frequency-hop receiver*. It is specifically designed for the type of interference that will be present in frequency-hop radio networks [10-12, 26], and it can account for the net effects of partial-band jamming, multiple-access interference, and narrowband interference without having to discriminate between these interference sources or respond to them individually.

Distributed routing schemes [1, 6, 19, 24, 27] can typically be characterized by four primary components: the metric used to assign weights or "distances" to individual links, the method used to propagate the link weights around the network, the algorithm used at the individual radios to determine the best paths based on these weights, and the procedures for storing and updating this routing information in the individual radios. The tier routing scheme, developed for the DARPA low-cost packet radio [6] and used in the SURAN protocol suite [1,9] is a distributed scheme which selects routes that minimize the distance measured in number of hops (a hop is a transmission from one radio to the next). The metric is very simple for tier routing: each good link is assigned weight 1 and each bad link is assigned weight ∞ . The information is propagated through the network by use of PROP's, and each radio transmits a PROP every 7.5 seconds. A Bellman-Ford type algorithm is used to determine the best routes based on the information available to

the radio from the PROP's that it has received from its neighbors, and the resulting routes are stored in the form of three tables: a neighbor table, a tier table, and a device table [5,6].

We take tier routing as a starting point, modify it, and add features of LRR in order to obtain more responsive adaptation in the presence of jamming and other interference. The primary difference between tier routing and LRR is the choice of metric. The metric for LRR is designed to reflect the way that interference will affect the performance of an FH radio receiver [12-14]. As a result, LRR achieves better adaptation for FH packet radio networks that must operate in the presence of partial-band interference and multiple-access interference.

The interference accounted for by the link metric for LRR is the interference *encountered* at the FH radio receiver. This interference may be a combination of multiple-access interference from other radios in the network, jamming, and other interference from outside the network (e.g., from friendly emitters operating in the same frequency band). The path or route metric used for LRR accounts for the interference that will be *encountered* by a packet traversing a particular route. This is in contrast to the schemes, such as least-interference routing [24], that account for the interference that a transmission will *cause* in the network. The interference that a receiver will experience is illustrated in Figure 4.1. In attempting to receive a packet from radio A, the receiver at radio B will experience jamming and multiple-access interference due to the transmission from radio C to radio D.




Figure 4.1

Interference encountered by the transmission from radio A

In the suggested implementation of LRR for FH packet radio networks, each radio forms an *interference statistic* for each of its frequency slots. This statistic may be based on the average energy during a certain period or on the demodulator outputs (or both). Let the vector of statistics be denoted by $\xi = (\xi_1, \xi_2, \dots, \xi_q)$, where ξ_i is the interference statistic for the i -th slot and q is the total number of frequency slots. The interference statistic for a given slot is a combination of all types of interference, so it accounts for partial-band jamming, interference from other FH transmissions, and narrowband interference. In this version of LRR, there is no intent to distinguish between these various types of interference. The vector ξ is the basis for one component of the resistance of the links into radio B. For omnidirectional antennas, this component will be the same for all of B's incoming links. Next, consider a function f that maps the interference statistic ξ_i to an interference level $f(\xi_i)$ for each frequency slot. The most elementary function suitable for this purpose is binary, and one such binary function gives a simple good vs. bad classification: $f(\xi_i) = 0$ if $\xi_i \leq \Lambda$ and $f(\xi_i) = 1$ if $\xi_i > \Lambda$. We can combine these statistics into a single interference measure for radio B in a number of ways. For instance, we can simply average the interference levels for the q frequency slots, in which case the *interference factor* for radio B is

$$I(B) = q^{-1} \sum_{i=1}^q f(\xi_i)$$

Notice that for the simple binary function f described above, the interference factor is just the ratio of the number of frequency slots with interference to the total number of frequency slots. If we are unable to monitor all frequency slots, estimates of $I(B)$ can be based on the fraction of the slots that we can monitor.

In more sophisticated versions of LRR, other information can be utilized. This other information might include the number of packets in radio B's buffer, the recent history of traffic near radio B, and an erroneous symbol count from

previously received packets. Such information can also be converted into an appropriate interference component and combined with the interference factor using suitable weightings. For example, for each neighbor, the radio should keep track of the number of symbols received in error in packets that were transmitted to it from that neighbor.

If $E(A,B)$ denotes the factor that accounts for the error rate experienced in recent transmissions from A to B, this factor can be combined with the interference factor to give the *resistance* for the link from A to B:

$$W(A,B) = \alpha I(B) + \beta E(A,B),$$

where α and β are coefficients that can be selected to give the desired emphasis on the interference factor relative the error-rate factor. Larger values for the resistance $W(A,B)$ correspond to a poorer quality link from A to B. The *resistance* for a particular path through the network is the sum of the metrics for the individual links. For instance, the path A-B-C-D has resistance

$$R(A,B,C,D) = W(A,B) + W(B,C) + W(C,D).$$

The goal is to select routes that correspond to the smallest value of the resistance.

For our simulation, the initial implementation of LRR assumed the number of frequency slots jammed and the number of interfering transmissions at a radio could be determined and the sum of these was taken as the resistance. In practice, it will be difficult to distinguish between these two types of interference.

Furthermore, the sum of these two statistics does not give a very good indication of the success probability on the links. In fact, the dependence of the probability of error on the number of interfering transmissions is not linear.

It is more realistic to assume that the radio will be able to estimate the number of frequency slots that have interference. We updated LRR to use the number of frequency slots with interference to estimate the conditional probability that a packet can be acquired and decoded given that the radio is not busy with

another transmission. The resistance is then taken to be the negative of the logarithm of this probability. This makes the resistance a function that increases in the appropriate way as the probability of success decreases. As before, the resistance for a path is just the sum of the resistances of each link. Equivalently, the path resistance is the negative of the logarithm of the product of the success probabilities for each link of a path. If successes were independent for each link (which they are not), the product of the success probabilities is the success probability for the path; nevertheless, this product is a reasonable estimate of the success probability for the path, at least for purposes of route determination. So, the resistance for a path can be viewed as the negative of the logarithm of a simple estimate of the end-to-end success probability for that path.

Because our simulation assumes the frequency hopping is asynchronous, the expected number of frequency slots that contain some type of interference is given by

$$N = q\{1 - (1-p)(1-2q^{-1})^J\}$$

where p is the fraction of frequency slots that are jammed, q is the total number of frequency slots, and J is the number of transmissions from other frequency-hop radios that are within range. If N were known, then Nq^{-1} could be used as an estimate for the probability that a given dwell interval has interference, and this estimate could be used to calculate an estimate of that receiver's success probability (i.e., the probability that a packet can be acquired and decoded at the radio in question).

The types of interference included in the simulation are multiple-access interference and partial-band jamming. We assume the radio is able to estimate N , but not necessarily able to determine the number of interfering frequency-hop signals or the fraction of the band that is occupied by the partial-band jammer. This is consistent with implementation constraints in current and next-generation

frequency-hop radios. In practice, N would have to be estimated by measuring the average number of frequency slots with interference during the packet interval or counting errors at the decoder output for previous packet receptions. Test symbols can also be employed to obtain an estimate of N .

5. Simulation Development

During Phase I, we have greatly enhanced and extended our simulation capabilities. In particular, a general framework has been included for adaptive protocols, mobile partial-band jamming has been added, and our new adaptive routing protocol, least-resistance routing (LRR), has been implemented. In this section, we describe some of the approaches followed in the development of our FH packet radio network simulation.

5.1 Simulation of Adaptive Routing

In our simulation, adaptive routing is carried out in a decentralized, distributed fashion. The result is a primary and secondary route from each radio to each other radio in the network. Instead of storing the entire route to each destination, a radio stores only the first link in this route, which is called the *outgoing* link for that route. This is a link from the radio to the neighbor that represents the next radio in the route. Upon receipt of a packet over this link, the neighbor forwards the packet according to its routing table, and the process continues until the packet either reaches its destination or is discarded because of an excessive number of transmission failures.

In general, routing algorithms have several different components that account for various features of particular networks. During the Phase I investigation, the primary interest has been in adaptive routing in the presence of jamming. Therefore, it has not been necessary to implement all features of the

routing algorithms that we have considered. In particular, we have not implemented all of the congestion control features of LRR.

There have been three major steps in the enhancement of our simulation program to permit the study of LRR and permit comparisons of LRR with other routing methods. First, the SURAN tier routing algorithm [5,6] was implemented to provide the basic backbone for LRR and other adaptive routing techniques. Second, the LRR algorithm was implemented by extending some of the procedures used in tier routing and by introducing new techniques to determine the interference environment throughout the network. Third, as simulation results were obtained, some observations were made for algorithms to improve the performance of LRR, and many of these improvements were implemented.

5.2 Implementation of Modified SURAN Protocols

Some modifications to the SURAN protocols are required to make them appropriate for use in a network with receiver-directed, frequency-hop transmission. This is the case for the forwarding protocols, for example. In a receiver-directed network, transmissions are made on the hopping pattern of the intended receiver, so only that receiver will attempt to acquire the transmission. The SURAN alt-routing procedure [5,6] relies on all radios being able to receive all transmissions that they can hear. In the context of frequency-hop spread spectrum, this is satisfied, for example, if all transmissions take place on the same hopping pattern, a restriction that we prefer not be imposed in general. To facilitate the use alternative routes in a network with receiver-directed spread-spectrum transmissions, each radio has a primary and secondary outgoing link for each destination. After a certain number transmission attempts are unsuccessful on the primary link, the packet is forwarded on the secondary link.

The SURAN protocols require that a routing table entry that uses a bad link be marked as bad and remain so for a certain period of time known as the *hold-down* time. This period is selected to provide time for the transmission of a certain number of PROP's in order to spread the information that the route is bad. The use of this hold-down time also helps prevent the formation of route loops. We elected not to use a hold-down time for two reasons. First, in our protocols, each radio stores a secondary outgoing link for each destination, and the forwarding protocol permits it to switch to this secondary link when too many failures have occurred on the primary link. Second, we included additional tests to be performed during routing table updates, and these reduce the probability of the formation of a route loop.

Inclusion of tier routing in our simulation was accomplished in three stages. First, the protocols necessary for the transmission of PROP's were implemented. In the second stage, neighbor tables were added. In the final stage, the routing tables were included. The approach followed in each of these three stages is described in the paragraphs that follow.

In our present simulation, PROP transmissions occur on a common hopping pattern, so all radios within range of a PROP transmission can potentially receive it. Each radio transmits a PROP every x time intervals (e.g., $x = 240$ in some cases). This means that a PROP is being transmitted somewhere in the subnetwork approximately once every $x/8$ time intervals for the subnetwork with eight radios (or $x/12$ for the subnetwork with twelve radios). Once a PROP has been generated at a radio, it is scheduled for transmission in the first time interval that was not previously scheduled for an acknowledgement transmission. Because a radio does not attempt to predict when the PROP will be transmitted by neighbor radios, PROP transmissions may occur when one or more radios are busy transmitting. For this and other reasons, not all PROP transmissions will be received. Although during

the early work in Phase I we have assumed the PROP transmissions are not affected by the jammer or other interference, this assumption is being removed in our later work in Phase I.

A radio begins transmitting a PROP in one of the first three dwell interval of its scheduled time interval; the dwell interval is chosen randomly with each interval being equally likely. Start times for other types of transmissions are chosen at random from the fourth through eighth dwell intervals of a time interval. Thus, in each packet interval, a radio that is in the receive mode will first look for the beginning of a PROP transmission on the common hopping pattern in the first three dwell intervals; if none is detected, the radio will then look for other transmissions on its own hopping pattern in the other five dwell intervals. This approach gives PROP's higher priority than other types of transmissions. As a result of having to look for incoming packets on two different hopping patterns, the time required for the frequency-hop radio to switch hopping patterns will be an important parameter in the eventual implementation of this scheme in an actual FH packet radio network.

For the next step in the inclusion of tier routing, we implemented the *neighbor table*, which is the table that is maintained at each radio to store statistics concerning the numbers of transmissions made and the number of packets received. Because of jamming, multiple-access interference, and conflicts between transmissions and receptions, the number of transmissions from radio A to radio B will exceed the number packets that B actually receives from A. The good neighbor table at radio A includes data on the number of transmissions A has made to each neighbor, the number of packets A has received from each neighbor since the last PROP was received from that neighbor, and the number of transmissions that each neighbor reports it has made to A prior to the last PROP from that neighbor. These data are used to classify the neighbors of radio A as either good or bad.

A routing table is maintained at each radio that lists information on the routes to each destination in the network. Information on the primary route to a given destination includes a specification of the primary outgoing link and the distance (i.e., number of hops) to that destination via the primary route. Similar information is given for the secondary route. Only good links are used for the routes in the routing table.

Routing information is passed around the network via PROP's: each PROP from a given radio identifies the primary outgoing link for each destination and the distance to each destination via the primary route. When a PROP is received from a good neighbor, the routing tables are updated to reflect the changes in the primary routing table of that neighbor. In addition, when a packet is received with a destination that no longer has a routing table entry, a flag is included in the acknowledgement so that the incorrect routing table entry will be cleared at the radio that receives the acknowledgement.

5.3 Implementation of LRR for FH Packet Radio Networks with Partial-Band Jamming

During Phase I, we completed the implementation of the simulation of the basic LRR method and some extensions. This was accomplished by extending some of the SURAN protocols and by adding some new protocols. In LRR a packet is forwarded on the outgoing link that currently has the least resistance to the destination. The distance information from the neighbor table is used to provide a controlled method for updating the resistances. This helps avoid route loops and other problems.

In the LRR approach, each radio obtains a quantitative measure, according to a prescribed metric, of the interference in its neighborhood (see Section 4). This interference contributes to the resistance of each route that passes through that radio. The interference information for radio B is determined during each time

interval that radio B is not transmitting. For the first metric for the Phase I simulations, the interference at radio B is the sum of the number of transmissions that are within range of radio B and the number of frequency slots that are jammed at radio B. The *interference factor* for radio B is the interference at radio B divided by the total number of frequency slots. In general, we suggest that the resistance for the link from radio A to radio B be a weighted sum of the interference factor at radio B and other measures of link quality and congestion, but, as a first step in implementing LRR, we used the interference factor only. For this measure, all links into radio B have the same resistance. This metric assumes that the fraction of the band jammed and the number of interfering transmissions are known precisely by each radio. We have also investigated the sensitivity of the network performance to this assumption by examining the effects of errors in the measurements of the number of frequency slots jammed. The second metric we have investigated assumes the radios can determine the number of frequency slots that have some type of interference. From this the probability a packet can be acquired and decoded is estimated and the resistance is taken to be the negative logarithm of this probability. These metrics are discussed in greater detail in Section 4.

The resistance for each route is updated by three methods. In the description of these methods, it is helpful to consider a particular route connecting radio A to radio D, namely A-B-C-D. The first mechanism permits each radio to learn of changes in its neighbor's interference factor. This is accomplished by including the interference factor for a radio in the header of each packet transmission and acknowledgement transmission made by that radio. Thus, when radio B receives one of these types of transmissions from radio C, radio B knows the current value for C's interference factor, and radio B updates the resistance values for all routes that go through radio C, including the route from B to D.

The second mechanism permits the information about neighbors to be passed on to other radios along the route. Suppose a packet destined for radio D is transmitted by radio A over the link to radio B. When B acknowledges this packet, B's current value for total resistance of the route from B to D is given in the header of the acknowledgement packet. This information and the interference factor at B (which radio A knows from the first mechanism) completely specify the current resistance for the route from A to D, and radio A updates its routing table accordingly. As packets are forwarded along the route from A to D, changes in resistance values along the route will be reflected back to A in this manner.

Finally, resistances are also updated by PROP transmissions using a Bellman-Ford type algorithm. A PROP sent by radio A includes the A's interference factor and the total resistance from A to each destination. The routing tables are ordered so that, for each destination, the primary table lists the route with the smallest resistance. If a particular destination cannot be reached, the resistance from A to that destination is set to infinity.

In the development of the LRR algorithm, we found it beneficial to include some of the techniques of tier routing to help avoid generating routes that contain loops. Loops can arise during the execution of the Bellman-Ford algorithm, but, when the algorithm converges, the route between with the least-resistance will not contain any loops. Unfortunately, since the resistances are changing continually, the Bellman-Ford algorithm may never converge. To reduce the number of route loops that are formed, some of the methods of SURAN's tier routing are incorporated into the routing algorithm. Other modifications are also made in order to obtain both a primary and a secondary route for each source destination pair.

A radio's routing table contains, for each destination, not only the resistance to that destination but also the distance (i.e., number of hops) to that destination.

The distance and the resistance from a radio to each destination are both specified in PROP's transmitted by that radio. When a PROP is received, the routing algorithm examines the change in the length of each route for evidence of a loop. Also, when a packet is forwarded, the distance to the destination is included in the packet's header so that the progress of the packet can be monitored to determine if the packet is traveling in a loop. If a loop is detected, a flag is included in the packet's acknowledgement to indicate that presence of this route loop.

There is also a one-bit flag in the header of each packet that is used to indicate lateral forwarding has taken place. The term, *lateral forwarding*, which is actually a contradiction in terminology, is used to describe the successful transmission of a packet in a way that does not move the packet any closer (in terms of the number of hops) to its destination [5,6]. This may happen, for example, when the initial forwarding attempts on the primary route have been unsuccessful, and the final attempts are made on a longer secondary route. Of course, the resistance to the destination can decrease in lateral forwarding, even if the distance stays the same. The flag is turned on when a packet has been forwarded laterally, and this prevents the packet from being forwarded laterally again at the next radio. Once progress toward the destination has been made by the packet, the flag is turned off and the packet can be forwarded laterally if necessary.

In preparation for developing new forwarding algorithms, we examined more closely the behavior of the routing algorithm. We found that few routes form which contain loops, and seldom is a radio without a route to each destination, but often the primary route is not the shortest-length path that is available. Also, we have observed that the resistance due to a partial-band jammer that occupies 40% or more of the band is always much greater than the resistance for any route through the subnetwork that does not pass through a jammed radio. For a single jammer, there are multiple routes between each pair of radios in our subnetwork that do not

pass through the jammed radio, so routes that pass through the jammed radio are only rarely stored in a routing table. After the jammer moves, the resistance for the routes through the new jammed radio are updated quickly, but the old routes through the previously jammed radio are not discovered as quickly. These observations have prompted two new protocols for the simulation program. First, we developed the *n*-packet forwarding protocol which proportions traffic between the primary and secondary links based on the difference in resistance between the paths. The second protocol generates extra PROP transmissions when the interference environment at a radio changes significantly. These extra PROP's are referred to as *event-driven* PROP's (as opposed to the PROP's a radio transmits at regular intervals).

The *n*-packet forwarding protocol begins by assigning the first $n-1$ packets for a specific destination to the primary outgoing link for that destination, and the n th packet for that destination is assigned to the secondary outgoing link. Forwarding continues in this manner with every n th packet for the same destination being assigned to the secondary outgoing link. The value of n depends on the difference in resistance between the primary and secondary routes to that destination. If the primary outgoing link for that destination changes, the algorithm is restarted at the time of the change. For any particular packet, the first three forwarding attempts are on the outgoing link assigned by the *n*-packet forwarding protocol. If all three fail, the packet is assigned to a different link for the remaining three forwarding attempts.

Recall that the resistance for a path is the negative of the logarithm of the estimated probability of success for a packet routed along that path. Let p_p and p_s denote the probability of success for the primary and secondary routes respectively, and let $\Delta = p_p - p_s$; note that $p_p \geq p_s$. Table 1 shows two rules for choosing n based on Δ .

TABLE 1. Threshold for the n -packet forwarding protocol

n	n -packet I	n -packet II
2	$\Delta \leq 0.1$	$\Delta \leq 0.01$
4	$0.1 < \Delta \leq 0.25$	$0.01 < \Delta \leq 0.3$
7	$0.25 < \Delta \leq 0.5$	$0.3 < \Delta \leq 0.6$
∞	$0.5 < \Delta$	$0.6 < \Delta$

The introduction of the event-driven PROP transmissions was prompted by the observation that after certain changes in the network, the update method took too long to locate new low-resistance routes. Without the event-driven PROP's, when the jammer moves, the resistance for the routes that pass through the newly jammed radio are quickly updated, but the routes through the radio that is no longer jammed are not found and added to the routing table very quickly. The event-driven PROP protocol is designed to help detect the routes that have become available because the jammer has moved. Only a radio that detects a change in its environment generates such a PROP, and a radio that receives one of these PROP's does not create a PROP in response.

In the process of developing this protocol some additional changes were made. A smoothing function was added for each radio to keep track of its local resistance. Each radio keeps the last three measurements of the estimated probability of success for a path. After a new resistance measurement is made, the radio compares the new probability of success with its oldest stored probability of success. When the probability of success has decreased by more than 30% (i.e., the resistance has increased) the radio chooses at random one of the next five time intervals for a PROP. When the probability of success has increased by more than 30% the radio waits five time intervals and chooses one of the following five time intervals for a PROP; again each time interval is equally likely. Delaying the PROP

in the second case allows the bad news to propagate first, increasing the probability that the neighboring radios will remove routes through the jammed radio and add routes through the radio that is no longer jammed. Requiring the resistance to change by at least 30% ensures that almost all the PROP's generated are caused by the jammer moving in or out of range of a particular radio (for the jamming levels considered in this report).

5.4 Models for Jamming and Multiple-Access Interference

In a FH packet radio network with partial-band jamming, the probability of error at a given radio receiver is a function of the fraction of the band jammed, referred to as the jamming level, and the total number of FH transmissions that are within range of that receiver. The jamming level and number of FH transmissions make up the interference environment for the radio. Together, these determine the probability that a particular dwell interval will have interference. In our model of the effects of partial-band jamming, it is assumed that all symbols transmitted in dwell intervals that are hit are erased at the receiver. This is a very good assumption for high-power partial band jamming if the radios develop side information [12] and uses it to erase unreliable symbols. However, an investigation of the effects of imperfect side information should be carried out in the future.

The present simulation allows for a single stationary or mobile partial-band jammer, and the number of FH transmissions is generated by a Markov chain model. The jammer is either in a fixed location near radio 1, or else it moves periodically between radios 1 and 3. In the present version of the simulation, only *spot jamming* is considered; that is, at any given time, the jammer effects transmissions made to one radio only.

We consider two different interference environments. The first is characterized by a high jamming level and a low level of FH transmission

interference; the second has a medium jamming level and a medium level of FH transmission interference. Based on our analysis of packet error probabilities, we decided that the high jamming level should correspond to partial-band jamming that occupies 55% of the band. The medium jamming level corresponds to 40% of the band being jammed. The FH transmission interference consists of transmissions from FH radios that are part of the subnetwork and transmissions from radios that are not part of the subnetwork (e.g., the unnumbered radios in Figure 3.1). In our simulation, the subnetwork is divided into five regions in order to keep track of the number of FH transmissions from outside the subnetwork. Four of the local regions are shown in Figure 5.1; a global region, which is not shown, includes all radios in the subnetwork.




Figure 5.1 Interference regions for the subnetwork

A discrete-time Markov chain for a given region specifies the number of interfering transmissions heard by radios in that region due to transmissions from

outside the subnetwork. Each Markov chain is independent of all others. The Markov chain for the global region is identical for both interference environments, and it is shown in Figure 5.2.




Figure 5.2

Markov chain for the all regions of the first interference environment and for the global region of the second interference environment




Figure 5.3

Markov chain for the local regions in the second interference environment

For the first interference environment, the Markov chains for the local regions are the same as for the global region. For the second interference environment, the state space and transition probabilities for Markov chains for the local regions are shown in Figure 5.3. Notice that this latter model corresponds to a much higher level of FH transmission interference than the model of Figure 5.2. During Phase I, we have been interested primarily in the behavior of LRR in the presence of jamming, and we have not included the aspects of LRR that can deal with traffic congestion. Because of this, we feel that the best test of the capabilities of LRR is achieved by use of the first interference environment (i.e., high jamming level, low level of FH transmission interference).

We assume that a given packet will not be received if the intended recipient is busy. A radio is busy if it is in the transmit mode, if it acquires a different packet first, or if another transmission arrives in the same dwell interval as the packet in question. Even if the intended recipient is not busy, the packet will not be received if the receiver cannot acquire synchronism due to a hit in the synchronization intervals, or if the radio cannot decode the packet because of an excessive number of hits. The conditional probability that the packet will not be received given that the intended recipient is not busy is determined by the interference environment. Throughout the simulation, the (32,16) Reed-Solomon code is employed, and the number of frequency slots is 200.

6. Simulation Results for Least-Resistance Routing

We have obtained four sets of simulation results in our study of the performance of LRR in a FH network with partial-band jamming and interference due to other FH transmissions. The first set of simulations was designed to permit comparison of the performance of LRR and tier routing in a FH network with a stationary jammer. The second set simulates LRR and tier routing in a FH network with a mobile jammer. The effects of errors in the resistance values were examined in the third set of simulations. The fourth set considers some improvements to the basic LRR algorithms and simulation model. Each of the first two and fourth set of simulations was run for the two interference environments described in Section 5.4: (1) a high jamming level and a low level of FH transmission interference and (2) a medium jamming level and medium level of FH transmission interference. The first interference environment corresponds to partial-band jamming with jamming present in 55% of the band; in the second, only 40% of the band is jammed. The FH transmission interference models for these two environments are depicted in Figures 5.2 and 5.3.

The performance of the routing protocols is determined from the throughput and end-to-end success probability obtained from the simulations. Certain packets are marked and these are used to obtain the necessary statistics. The radios are numbered as in Figure 3.1, and the packets that are marked are those sent on the routes with source-destination pairs (5,6), (6,5), (7,8), and (8,7). The simulation data obtained from the marked packets are combined to give a single value for each of the performance measures. One reason for tracking the marked packets only is that, as can be seen in Figure 3.1, these packets must travel at least three hops. Packets that only travel one or two hops are not influenced by the routing algorithms as much as packets that travel farther.

The *end-to-end throughput* is defined to be the average number of marked packets that reach their destinations per time interval. The *end-to-end probability of success* is defined to be the average fraction of marked packets that reach their destinations.

6.1 Stationary Jammer

The first set of simulations was run for both interference environments for the eight radio subnetwork. The jammer is stationary and it interferes with radio 1 only. Figures 6.1 and 6.2 show the performance when 55% of the band is jammed, and Figures 6.3 and 6.4 give the corresponding results when 40% of the band is jammed. These results show that the performance of tier routing is sensitive to the threshold that is used in determination of the link status, and only when the threshold is properly chosen does tier routing perform as well as LRR. When the threshold is too high (e.g., 0.4), there are too many false alarms (i.e., good links that are declared bad). Tier routing performs much worse than LRR when the threshold is too high.

When the threshold is low (e.g., 0.2), there are many misses (bad links that are declared good), but the performance of tier routing is about the same as LRR. For low thresholds, the stationary jammer location is easily identified (at either jamming level), a route around the jammer is found for the primary table, and rarely is this route ever removed from the primary table. Also, because secondary forwarding is used, if the primary link is to a jammed radio, the secondary link will be to a radio that is not jammed.

In fact, when 55% of the band is jammed and the tier routing threshold is low, tier routing performs slightly better than LRR at the 0.04 packet generation probability. Both of these routing algorithms can easily adjust the routing tables to avoid the jammed radio. The routing tables stabilize when tier routing is used unless there are many false alarms; there are not many misses at the high jamming level. Tier routing always lists the shortest length path with no bad links in the primary routing table. On the other hand, when LRR is used, the routing tables continually change as the resistances change at the other radios. As these resistances change, LRR may occasionally use a longer route that has smaller resistance. Only limited conclusions about tier routing can be drawn from the simulations with a stationary jammer, because many of the hardest problems in establishing routing tables do not arise in the stationary jamming environment. In particular, the time required to determine that a link is bad, which is rather long for tier routing, is not important for stationary jamming.

6.2 Mobile Jammer

In the second set of simulations, the jammer moves between radios 1 and 3. First consider the results for the eight radio subnetwork. We consider two different periods for the jammer's movement. The *period* for the mobile jammer is the number of intervals between position changes; for example, if the period is 1000, the

jammer moves from a radio 1000 time intervals after it arrives at that radio. We considered jammer periods of 1000 and 2000 in the simulations.

The PROP interval is the number of time intervals between transmissions of a PROP from a given radio. A PROP interval of 240 was employed in simulations for both jammer periods, and a PROP interval of 160 was also investigated for a jammer period of 1000. Both interference environments were simulated for a jammer period of 1000 and PROP interval of 240, but only the first interference environment (high level of jamming, low level of FH transmission interference) was examined for a jammer period of 2000 or a PROP interval of 160.

The results of the simulations for a jammer period of 1000 and the first interference environment (55% of the band jammed) are shown in Figures 6.5 and 6.6. The performance advantage of LRR over tier routing is clearly seen in Figure 6.5. A number of factors contribute to the improved performance of LRR. The LRR algorithm is not limited by having to wait for a PROP to update routing tables, because some of the resistances are updated after every transmission. Also, the LRR algorithm does not require a significant number of transmissions in order to determine the status of a neighbor. Instead, the resistance at a radio is measured during every time interval the radio is not transmitting. The tier routing algorithm does not react as quickly to a jammer, because it must wait for at least one PROP after the jammer has moved to detect any change in the conditions along the routes.

The version of LRR used for these simulations does not provide a significant improvement over tier routing when the jamming level is low, but this is as expected. The present version of LRR tends to avoid use of the jammed radio altogether which greatly reduces the number of routes available. Also, because the connectivity is so limited in the eight radio subnetwork, there are no routes that are longer than the shortest length route that can provide a better path.

Figures 6.7 and 6.8 show the change in network performance when the jammer moves more slowly (a jammer period of 2000). The graphs show the sensitivity of tier routing to the frequency with which it must adjust to new interference patterns. LRR is much more resilient to the different jamming scenarios.

Next, consider LRR and tier routing in the twelve radio subnetwork as shown in Figures 6.9–12. The jammer period is 1000 in these simulations, and both jamming levels are considered. In this subnetwork LRR performs much better than tier routing even when the threshold for tier routing is adjusted to give the best performance. The large network offers more alternative routes to LRR that are longer than the shortest length route but have less resistance. In the small subnetwork tier routing performed its best when the threshold was 0.2 while in the larger network it performs the best when the threshold is 0.4. Since the connectivity of the smaller network was limited, tier routing performed the best when few links were declared bad. When more alternative routes are available, the performance improves by declaring more links bad and using the alternative routes.

6.3 Effects of Errors in Resistance Values

So far we have assumed that each radio can determine the number of interfering transmissions and the number of slots that are jammed. To simulate the effects of errors in measuring these quantities, the resistance values are randomly modified. When the simulation calculates the resistance at a particular radio, it decreases the value by $x\%$ with probability $1/3$, increases the value by $x\%$ with probability $1/3$, and leaves the value unchanged with probability $1/3$. Percentages of 25% and 50% are considered for x . The errors are included in the simulation runs with a jammer period of 1000 and 55% of the band jammed.

The results are shown in Figures 6.13 and 6.14. There is almost no difference in performance between the three simulations. This is very encouraging, but some of this is probably due to the topology. Even if a route is selected that contains a jammed radio, a secondary route exists at each radio that can avoid the jammer. Also, since the fraction of the band jammed is large, the resistance for the jammed radio is always much greater than for the other radios. Further investigations of the effects of errors in the measurements of the interference will be conducted in the future.

6.4 Additional Protocols for LRR

In the final set of simulations only the twelve radio subnetwork is used, and both interference environments are used. The jammer period is 1000 time intervals, and each radio schedules a PROP every 240 time intervals.

Under the first interference environment (55% of the band jammed), the new LRR metric (based on the link success probabilities) results in a significant improvement over the old metric (based on the number of interfering transmissions), as shown in Figures 6.15 and 6.16. The new metric provides a better indication of the resistance for the routes, and when a longer route should be used instead of the shortest-length route. With 40% of the band jammed, the difference between the two metrics is not as large as with the first interference environment, as shown in Figures 6.17 and 6.18. The difference between the performance of the two interference environments is linked to the higher level of FH transmission interference in the second interference environment. This model changes the FH transmission interference over a large range, making it more difficult to keep accurate route resistances. One of the issues to address in future work is how to handle changes in the interference environment that cannot be followed with the routing algorithm. We have addressed this problem in part by examining some

forwarding algorithms; forwarding algorithms can adapt to local fluctuations in the environment quicker than routing tables can be updated.

Next we considered the n -packets forwarding protocol (described in Section 5.3) in the twelve radio network (all the remaining simulations use the new metric for LRR only). The n -packets forwarding protocol slightly reduced the performance in both interference environments compared to LRR with the primary 6/3 forwarding protocol as shown in Figures 6.19–22. The motivation was to force some packets on the secondary routes to increase the feedback about the neighbors of a radio. Recall that the resistance at a radio is included in each packet and acknowledgement, and the route resistance to the destination for a particular packet is included in its acknowledgement. One problem may be that occasionally a route will be included in the routing table that is a few hops longer than the best route possible. This is more likely to occur for the secondary route, and the secondary route is updated less often since it is not used as often.

Next, we add the event-driven PROP protocol and included the effects of interference for the PROP transmissions. We also considered the case when the time between regular PROP transmissions is doubled to 480 time intervals. Figures 6.23–26 show the performance for both interference environments. The performance is not effected very much by any of these additions. The performance increases slightly at the higher packet generation probabilities when event-driven PROP's are used, the PROP's are subject to interference, and fewer PROP's are scheduled. This suggest that fewer PROP's are need and some throughput may be lost because of excessive PROP traffic. Also, any improvements from using event-driven PROP's may not be very noticeable since they are generated only when the jammer moves (i.e., once every 1000 time intervals).

While using the event-driven PROP's we examined again some of the forwarding protocols. With 55% of the band jammed, LRR with the primary 6/3

and 6/6 forwarding protocols perform nearly the same, as shown in Figures 6.27–28; LRR is able to determine the best link for the primary routing table and the secondary route is not needed for forwarding packets (the secondary route is still used when the primary route's resistance increases). Also, there is not much extra FH transmission interference so information about the resistance for a link does not change too often. The n -packet II forwarding protocol does not work as well; the best link is the primary link so it does not help to use the secondary link. When 40% of the band is jammed the primary 6/6 forwarding protocol performs worse than the primary 6/3 and n -packets forwarding protocols, as shown in Figures 6.29–30. The background interference changes much more in the second interference scenario so it is beneficial to try the secondary link when the primary link has been unsuccessful. Also, with the interference levels changing more, the n -packets protocol is able to update the resistance for more links because it forces some packets on the secondary links.

We examined the performance of LRR with event-driven PROP's and jamming of PROP transmissions when the packet generation probability is higher. As expected the performance decreases quickly as the network becomes more congested, as shown in Figures 6.31–34. We have not included any algorithms to control congestion. Ideally, as the generation probability increases, the probability of success will decrease, but, the throughput should not decrease as quickly.

Finally, we changed the model for generating packets at the marked and unmarked radios. The packet generation probability is reduced from $p/2$ to $p/10$ for the unmarked radios. Also p was increased so that the overall generation probability for the network is approximately the same as in the previous simulations (i.e., the same number of packets are in the system, but more of them are originating at the radios 5–8). The maximum throughput obtained is higher

than for the other packet generation model, as shown in Figures 6.35–38. But, the throughput is sharply reduced at the highest packet generation probability.




Figure 6.1. Throughput for stationary jamming with 55% of band jammed




Figure 6.2. Success probability for stationary jamming with 55% of band jammed




Figure 6.3. Throughput for stationary jamming with 40% of band jammed




Figure 6.4. Success probability for stationary jamming with 40% of band jammed




Figure 6.5. Throughput for mobile jamming with period 1000 and 55% of band jammed




Figure 6.6. Success probability for mobile jamming with period 1000 and 55% of band jammed




Figure 6.7. Throughput for mobile jamming with 55% of band jammed




Figure 6.8. Success probability for mobile jamming with 55% of band jammed




Figure 6.9. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.10. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.11. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.12. Success probability for network with 12 radios, mobile jamming with period 1000 and 40% of band jammed




Figure 6.13. Throughput for mobile jamming with period 1000 and 55% of band jammed




Figure 6.14. Success probability for mobile jamming with period 1000 and 55% of band jammed




Figure 6.15. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.16. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.17. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.18. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.19. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.20. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.21. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.22. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.23. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.24. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.25. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.26. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.27. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.28. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.29. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.30. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.31. Throughput for network for 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.32. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of band jammed




Figure 6.33. Throughput for network for 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.34. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of band jammed




Figure 6.35. Throughput for network with 12 radios, mobile jamming with period 1000, and 55% of the band jammed




Figure 6.36. Success probability for network with 12 radios, mobile jamming with period 1000, and 55% of the band jammed




Figure 6.37. Throughput for network with 12 radios, mobile jamming with period 1000, and 40% of the band jammed




Figure 6.38. Success probability for network with 12 radios, mobile jamming with period 1000, and 40% of the band jammed

7. Investigation of Pacing

As part of a subcontract to SIGCOM, ITT undertook an investigation of pacing using their network simulator [10, 26], which, for convenience, we refer to as the combat net radio simulator (CNRS). ITT's channel model is different from SIGCOM's in that all radios use a common FH hopping pattern. SIGCOM's simulation uses a common hopping pattern for the PROP transmissions only — all other transmissions use the hopping pattern of the intended receiver. As a result, CNRS can use passive or active acknowledgements.

Use of passive acknowledgements allows CNRS to use adaptive pacing. A radio monitors the time between when it forwards a packet and when it receives the acknowledgement for that packet to determine how long to wait before forwarding another packet. The detailed algorithm for this calculation is given in [26]. This calculation is not possible with active acknowledgements since the receiving radio sends an active acknowledgement to the sending radio immediately after receiving a packet from the sending radio; the delay between receiving a packet and sending an acknowledgement is constant instead of depending on the number of packets already in the transmission queue.


When the delay between forwarding a packet and receiving an acknowledgement is not available (when active acknowledgements are used) or is not used, constant pacing can be used. That is, when radio A has forwarded a packet to radio B and received an acknowledgement, radio A employs a fixed delay between the time it made the last transmission to B and when it makes the next transmission to radio B. ITT has explored constant pacing with passive acknowledgements only, and they have chosen to use three times the packet transmit time as the constant pacing delay. This coincides with the three frame nature of forwarding a packet: one frame for radio A to forward a packet to radio B, one frame for B to forward and acknowledge the packet, and one frame for B to receive an acknowledgement. Note,

SIGCOM uses constant pacing and active acknowledgements, and our constant pacing delay is four time intervals (one extra transmission is necessary because one transmission cannot both acknowledge and forward a packet).

ITT studied the network performance using CNRS with passive acknowledgements and both adaptive and constant pacing, and with active acknowledgements and no pacing delay. The eight-radio subnetwork consisting of radios 1-8 of Figure 5.1. Throughput and transmit time (i.e., end-to-end delay) results are shown for all packets and just the marked packets in Figures 7.1-6. At the lower packet generation rates all three protocols perform approximately the same. As the number of packets in the network increases the constant pacing protocol with passive acknowledgements performs the best; the simulations with the adaptive pacing protocol did not perform as well, and the simulations with active acknowledgements performed the worst. The adaptive pacing algorithm use parameters that were found to be optimal in an earlier experiment; an extensive study was not done to determine the best parameters for this network. Constant pacing has the advantage of not require the pacing parameters to be adjusted.

Note that only the packets that reach their destinations are counted in the transmit time. At the higher packet generation rates many of the packets are discarded before reaching their destinations. Such packets have to be detected by the end-to-end acknowledgement protocol, which is not included in this simulation, and these packets will have to be rescheduled for transmission at a later time.

SIMULATION RESULTS FOR AN 8 NODE NETWORK WITH CONSTANT PACING
 NETWORK LOADING - POISSON DISTRIBUTION



AVERAGE NETWORK LOADING (PACKETS PER SECOND)

Figure 7.1

SIMULATION RESULTS FOR AN 8 NODE NETWORK WITH ADAPTIVE PACING

PACING PARAMETERS $a = 0.95$ $C = 2$





Figure 7.2


SIMULATION RESULTS FOR AN 8 NODE NETWORK
PACKETS ACTIVELY ACKNOWLEDGED AT EACH NODE



AVERAGE NETWORK LOADING (PACKETS PER SECOND)

Figure 7.3

**SIMULATION RESULTS FOR AN B MODE NETWORK WITH CONSTANT PACING
NETWORK LOADING - POISSON DISTRIBUTION**



MARKED PACKET NETWORK LOADING (PACKETS PER SECOND)

Figure 7.4

SIMULATION RESULTS FOR AN 8 NODE NETWORK WITH ADAPTIVE PACING

PACING PARAMETERS $\alpha = 0.95$ $C = 2$





Figure 7.5

SIMULATION RESULTS FOR AN 8 NODE NETWORK
PACKETS ACTIVELY ACKNOWLEDGED AT EACH NODE



MARKED PACKET NETWORK LOADING (PACKETS PER SECOND)

Figure 7.6

8. References

- [1] D. Beyer, "Minutes of the SURAN working group meeting," Cambridge, MA, Oct. 1987; Menlo Park, CA, Feb. 1988; Cambridge MA, May 1988; and Richardson, TX, Oct. 1988.
- [2] N. Gower and J. Jubin, "Congestion control using pacing in a packet radio network," *1982 IEEE Military Commun.Conf. Record*, vol. 1, pp. 23.1.1-6, Oct. 1982.
- [3] B. Hajek, "Adaptive transmission strategies and routing in mobile radio networks," *Proc. 1983 Conf. on Inform. Sci. and Syst.* (Johns Hopkins Univ.), pp. 373-378, March 1983.
- [4] B. Hajek, "Adaptive packet radio networks," University of Illinois Technical Report, Oct. 1983 (reprinted in [19]).
- [5] J. Jubin, "Current packet radio network protocols," *IEEE INFOCOM Proc.*, pp. 86-92, March 1985.
- [6] J. Jubin and J. Tornow, "The DARPA packet radio network," *Proc. IEEE*, vol. 75, pp. 21-32, Jan. 1987.
- [7] T. Klein, "Research issues for Army communications," presented at the Army Research Office Long Term Strategy Workshop, May 1987.
- [8] G. Lauer et al., "Survivable protocols for large scale packet radio networks," *1985 IEEE Global Communications Conf. Record*, vol. 1, pp. 15.1.1-4, Nov. 1985.
- [9] B. M. Leiner, D. L. Nielson, and F. A. Tobagi, "Issues in packet radio network design," *Proc. IEEE*, vol. 75, pp. 6-20, Jan. 1987.
- [10] J. R. McChesney and M. B. Pursley, "Channel access/multiple access performance for SINCGARS packet switched radio network," *Proc. AFCEA-IEEE Technical Conference on Tactical Communications*, vol. 1, pp. 7-10, April 1986.
- [11] M. B. Pursley, "Adaptive spread spectrum radio networks," Report T-131, CSL, University of Illinois, Oct. 1983 (reprinted in [19]).
- [12] M. B. Pursley, "Frequency-hop transmission for satellite packet switching and terrestrial packet radio networks," *IEEE Trans. on Information Theory*, vol. IT-32, pp. 652-667, Sept. 1986.

- [13] M. B. Pursley, "Packet error probabilities in frequency-hop radio networks—Coping with statistical dependence and noisy side information," *IEEE Global Telecommunications Conference Record*, vol. 1, pp. 165-170, December 1986.
- [14] M. B. Pursley, "Tradeoffs between side information and code rate in slow-frequency-hop packet radio networks," *Conference Record, IEEE International Conference on Communications*, vol. 2, pp. 947-952, June 1987.
- [15] M. B. Pursley, "The role of spread spectrum in packet radio networks," *Proc. IEEE*, vol. 75, pp. 116-134, Jan. 1987.
- [16] M. B. Pursley and H. B. Russell, "Simulation of routing and pacing in a frequency-hop multiple-access packet radio network," presented at the May 1988 SURAN Working Group Meeting, Cambridge, MA, May 11, 1988.
- [17] M. B. Pursley and H. B. Russell, "Forwarding protocols in frequency-hop packet radio networks with partial-band jamming," presented at the October 1988 SURAN Working Group Meeting, Richardson, TX, Oct. 5, 1988.
- [18] M. B. Pursley and H. B. Russell, "Adaptive Forwarding and Routing in Frequency-Hop Spread-Spectrum Packet Radio Networks with Partial-Band Jamming," accepted for presentation at the 1989 Military Communications Conference, Boston, MA, Oct. 1989.
- [19] P. Sass (ed.), "Adaptive spread spectrum network research — ECCM and low probability of intercept for radio networks," Research and Dev. Tech. Rept. CECOM-TR-84-4, U. S. Army CECOM, April 1984.
- [20] N. Shacham, "Adaptive link level protocols for packet radio networks," *Proc. of INFOCOM '86*, March 1986.
- [21] N. Shacham and J. Westcott, "Packet radio networks: Future directions," *Proc. IEEE*, vol. 75, pp. 83-99, Jan. 1987.
- [22] A. Shohara, J. J. Garcia-Luna, M. G. Lewis, B. M. Sifford, J. R. McChesney, L. G. Matheson, and M. B. Pursley, "Design plan: SINCGARS packet switch overlay," prepared for the U. S. Army Communications- Electronics Command, Contract No. DAAB07-85-C-K581, May 1986.
- [23] J. Stevens, "SURAN network susceptibilities," SURAN Technical Note, SRNTN 39, Nov. 1985.
- [24] J. A. Stevens, "Least interference routing," *Proceedings of the 1988 Tactical Communications Conference*, Fort Wayne, Indiana, May 1988.

- [25] U.S. Army Electronics Coordinating Group, "U.S. Army Electronics Strategy Planning Workshop Report," 1987.
- [26] S. Vitalpur, "Protocol design and analysis of CNR packet radio networks," *Proceedings of the 1988 Tactical Communications Conference*, Fort Wayne, Indiana, May 1988.
- [27] J. Westcott and J. Jubin, "A distributed routing design for a broadcast environment," 1982 *IEEE Military Commun. Conf. Record*, vol. 3, pp. 10.4-1 through 10.4-5, Oct. 1982.